

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellants:	Oleg Kiselev		
Assignee:	Symantec Operating Corporation		
Title:	Remote Data Access For Local Operations		
Application No.:	10/722,701	Filed:	November 25, 2003
Examiner:	Chelcie L. Daye	Group Art Unit:	2161
Docket No.:	VRT0106US	Confirmation No.:	7436

Austin, Texas
July 14, 2008

Mail Stop: Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

SUBSTITUTE APPEAL BRIEF UNDER 37 CFR § 41.37

Dear Sir:

This brief is submitted in response to the Notification of Non-Compliant Appeal Brief dated June 13, 2008, and replaces the brief filed on June 6, 2008, in support of the Notice of Appeal filed February 6, 2008 by the appellants to the Board of Patent Appeals and Interferences from the Examiner's final rejection of claims 1-21. This brief includes a revised summary, which remedies the defects described in the Notification of Non-Compliant Appeal Brief. The fee for filing an appeal brief was paid with the brief filed on June 6, 2008. No new fee is required to be paid with this substitute brief.

I. REAL PARTY IN INTEREST

The real party in interest on this appeal is Symantec Operating Corporation.

II. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences related to this application.

III. STATUS OF CLAIMS

Claims 1-21 are pending in the application.

Claims 1-21 stand rejected.

The Appellants appeal the rejection of claims 1-21.

IV. STATUS OF AMENDMENTS

No amendments were filed subsequent to the final rejection of November 6, 2007.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Claim 1 describes a method that involves receiving (e.g., in operation 394 of FIG. 3G, described on lines 26-30 of page 23) a request to read a portion of data from first data storage (e.g., data storage 140A of FIG. 1). The request is received by a receiving module (e.g., storage manager / replicator 120A of FIG. 1) of a first host (e.g., primary node 110A of FIG. 1), e.g., as described on line 30 of page 24 through line 1 of page 25. The first host can access the first data storage, and the first host cannot access second data storage (e.g., data storage 140B of FIG. 1), e.g., as shown in FIG. 1 and described on page 10. The method also involves requesting (e.g., in operation 395 of FIG. 3G, described on line 30 of page 23 through line 2 of page 24) a requested portion of a copy of the data in the second data storage from a second host (e.g., node 110B of FIG. 1) that can access the second data storage.

The requested portion is received from the second host, and the portion of the data is read by reading the requested portion received from the second host (e.g., in operation 396 of FIG. 3G, described on lines 19-22 of page 24) and, when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage (e.g., in operation 397 of FIG. 3G, described on lines 22-25 of page 24).

Claim 13 describes a system that includes first receiving means (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) of a first host (e.g., primary node 110A of FIG. 1) for receiving a request to read a portion of data from first data storage (e.g., data storage 140A of FIG. 1), as described, for example, in lines 26-30 of page 23. The first host can access the first data storage, and the first host cannot access second data storage (e.g., data storage 140B of FIG. 1), as shown, for example, in FIG. 1.

The system also includes requesting means (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) for requesting a requested portion of a copy of the data in the second data storage (e.g., data storage 140B of FIG. 1) from a second host (e.g., node 110B of FIG. 1) that can access the second data storage, as described, for example, on line 30 of page 23 through line 2 of page 24. Second receiving means (e.g.,

storage manager / replicator 120A of primary node 110A of FIG. 1) for receiving the requested portion from the second host are also included.

The system further includes reading means (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) for reading the portion of the data by reading the requested portion received from the second host (e.g., by performing operation 396 of FIG. 3G, described on lines 19-22 of page 24) and, when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage (e.g., by performing operation 397 of FIG. 3G, described on lines 22-25 of page 24).

Claim 16 describes a system that includes a first receiving module (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) of a first host (e.g., primary node 110A of FIG. 1) to receive a request to read a portion of data from first data storage (e.g., data storage 140A of FIG. 1), as described, for example, in lines 26-30 of page 23. The first host can access the first data storage, and the first host cannot access second data storage (e.g., data storage 140B of FIG. 1), as shown, for example, in FIG. 1.

The system also includes a requesting module (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) to request a requested portion of a copy of the data in the second data storage (e.g., data storage 140B of FIG. 1) from a second host (e.g., node 110B of FIG. 1) that can access the second data storage, as described, for example, on line 30 of page 23 through line 2 of page 24. A second receiving module (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) to receive the requested portion from the second host is also included.

The system further includes a reading module (e.g., storage manager / replicator 120A of primary node 110A of FIG. 1) to read the portion of the data by reading the requested portion received from the second host (e.g., by performing operation 396 of FIG. 3G, described on lines 19-22 of page 24), and when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage (e.g., by performing operation 397 of FIG. 3G, described on lines 22-25 of page 24).

Claim 19 describes a computer-readable medium (e.g., one of elements 438, 442, 44, or 417 of FIG. 4) that includes first receiving instructions to receive (e.g., by

performing operation 394 of FIG. 3G, described on lines 26-30 of page 23) a request to read a portion of data from first data storage (e.g., data storage 140A of FIG. 1). A first host (e.g., node 110A of FIG. 1) can access the first data storage, and the first host cannot access second data storage (e.g., data storage 140B of FIG. 1). The computer readable medium also includes requesting instructions to request (e.g., by performing operation 395 of FIG. 3G, described on line 30 of page 23 through line 2 of page 24) a requested portion of a copy of the data in the second data storage from a second host (e.g., node 110B of FIG. 1) that can access the second data storage.

Second receiving instructions are included to receive the requested portion from the second host. Reading instructions read the portion of the data by reading the requested portion received from the second host (e.g., by performing operation 396 of FIG. 3G, described on lines 19-22 of page 24) and, when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage (e.g., by performing operation 397 of FIG. 3G, described on lines 22-25 of page 24).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

I. Whether claims 1-10 and 13-21 are patentable under 35 U.S.C. §103(a) over DeKoning et al. (USPN 6,691,245) (“DeKoning”) in view of Takeda (USPPN 2004/0172509) (“Takeda”).

II. Whether claims 11-12 are patentable under 35 U.S.C. §103(a) over DeKoning in view of Takeda and further in view of Carlson et al. (USPN 6,377,959) (“Carlson”).

VII. ARGUMENT

I. Claims 1-10 and 13-21 are patentable over DeKoning in view of Takeda

The cited art fails to teach or suggest reading a portion of data by reading the requested portion received from a second host and, “when a sub-portion of the portion of the data is not included in the requested portion [received from the second host, which was requested to provide the requested portion of a copy of the data in a second data storage], reading the sub-portion from the first data storage.” The Examiner acknowledges that this feature is not taught by DeKoning and relies solely upon paragraphs 61-65 of Takeda to teach this feature. Office Action mailed June 8, 2007 (“OA”), pages 3-4. The cited portions of Takeda recite:

[0061] Thereupon, the secondary host 100B then transmits a copy request ("journal copy request command") to the secondary disk array device 200B to initiate journal transfer from the primary disk array device 200A (Step 9300).

[0062] After receiving the journal copy request command, the secondary disk array device 200B issues a data read request to the primary disk array device 200A (Step 9310). The primary disk array device 200A transmits the requested data to the secondary disk array device 200B by executing the copy program 2110. Details of the journal copy processing will be described later.

[0063] On the other hand, the data that had been stored in PVOL 2212 before the journal acquisition process was started is not transferred to the secondary disk array device 200B even when the journal copy processing was started. Therefore, it is necessary to copy these data (hereafter "initial data") to SVOL 2214 from PVOL 2212. In the present embodiment, an initial copy process is used to transfer the initial data from the PVOL 2212 to SVOL 2214 (Step 9130). The initial data are transferred sequentially from the volume head area to the end of PVOL 2212 according to instructions of the host 100. This process may also be performed by allowing each disk array device 200 itself to execute the copy program 2110.

[0064] The initial copy and the journal copy processing may be performed asynchronously and in parallel. That is, the initial copy can be performed anytime after PVOL 2212 and SVOL 2214 have been specified based on the pair generation command, regardless of whether or not the journal acquisition process and the journal copy process has been performed or is being performed. However, as long as the initial copy has not been

completed, the contents of SVOL 2214 does not reflect PVOL 2212 even if the restoration process 9126 has been performed at the secondary disk array device 200B. The restore or restoration process involves updating or copying the data of PVOL 2212 in the SVOL 2214 using the journal that have been received from the primary disk array device 200A according to the copy process 9124.

[0065] In one implementation, the initial copy initiated by the secondary disk array device 200B issuing one or plural read commands to the primary disk array device 200A in order to reduce the load of the primary disk array device 200A.

Paragraphs 61-62 of Takeda describe how a second host (host 100B) initiates copying of data from the primary storage array to the secondary storage array using a journal copying process. Thus, this portion of Takeda simply describes copying data from the primary storage array to the secondary storage array.

Paragraphs 63-64 describe how some of the initial data on the primary storage array may not be captured by the journal copying process and is instead copied to the secondary storage array using an initial copy process. Thus, this portion of Takeda also describes copying data from the primary storage array to the secondary storage array. Paragraph 65 simply indicates that the secondary disk array device can initiate the initial copy process, just as the secondary disk array device can initiate the journal copying process.

Thus, in the cited portions of Takeda, all data that is copied or otherwise accessed is copied from the primary storage array to the secondary storage array. Takeda describes how all data will be obtained from the primary storage array using the combination of the journal copying process and the initial copy process and clearly does not disclose or suggest reading a sub-portion of data from the secondary storage array, if that sub-portion was not included in data received from the primary storage array (or vice versa). Thus, the cited portions of Takeda clearly neither teach nor suggest a scenario in which a sub-portion of data is not received from one data storage accessible via one host and is instead read from another data storage accessible via another host; instead, the cited portions of Takeda simply describe a scenario in which all data will be obtained from the same data storage.

In the Response to Arguments section of the Final Office Action mailed January 11, 2007 ("FOA1"), the Examiner states: "if the initial copy has not been completely

transferred (i.e., there is a portion of the data not transferred into the appropriate device) then the second device issues read commands to the primary device. This portion of Takeda corresponds to the claim limitations of a sub-portion of the requested portion is not available in a received portion and reading that unavailable portion from anything corresponding to the claimed first data storage.” “FOA1”, p. 11.

The Appellants first note that in the rejection of the claims, the Examiner has equated the secondary host of Takeda with the first host of claim 1, since the Examiner is stating that paragraphs 61-62 (which describe the journal copying process initiated by the secondary host) teach “reading the requested portion received from the second host.” In other words, since the cited portion of Takeda teaches the copying of data from the primary storage array to the secondary storage array, any data received is received from Takeda’s primary storage array, not from Takeda’s secondary host. For this to apply to the feature of claim 1 at issue, it appears that Takeda’s primary host is clearly being equated with the second host of claim 1. Furthermore, it appears that the primary storage array is being equated with the second data storage of claim 1, which is accessible by the second host of claim 1.

Accordingly, the fact that Takeda then teaches the secondary device issuing read commands to the primary device (as part of the initial copy described in paragraphs 63-65) simply means that Takeda is teaching requesting and then receiving more data from Takeda’s primary host. Since the primary host of Takeda is being equated with the second host of claim 1, this portion of Takeda, at best, teaches an action that would be equivalent to receiving both the requested portion and the sub-portion of the data from the second host of claim 1, and this scenario is clearly not what is described in claim 1. Thus, Takeda clearly does not teach or suggest “when a sub-portion of the portion of the data is not included in the requested portion [received from the second host, which was requested to provide the requested portion of a copy of the data in a second data storage], reading the sub-portion from the first data storage.”

Stated another way, the rejection attempts to equate the primary host of Takeda with both the first host of claim 1 (when expressing how the cited art allegedly teaches “reading the sub-portion from the first data storage”) and the second host of claim 1 (when expressing how the cited art allegedly teaches “reading the requested portion

received from the second host”). This application of Takeda is clearly inconsistent with the express terms of claim 1, which clearly recite two hosts and expressly recite the actions performed by each host.

In response to the above arguments, the Examiner again equates the journal copying process, which reads data from Takeda’s primary volume in order to copy that data to Takeda’s secondary volume, with claim 1’s act of “reading the sub-portion from the first data storage.” OA, pages 11-12. This interpretation is inconsistent with the positions asserted elsewhere in the rejection, where the Examiner has equated Takeda’s secondary volume with the first data storage. See, e.g., OA, page 4, lines 1-2.

Furthermore, merely copying data from Takeda’s primary volume to Takeda’s secondary volume does not teach or suggest reading a portion of data by reading the requested portion received from a second host and, “when a sub-portion of the portion of the data is not included in the requested portion [received from the second host, which was requested to provide the requested portion of a copy of the data in a second data storage], reading the sub-portion from the first data storage.” The cited portions of Takeda do not teach or suggest reading a sub-portion of requested data from one storage volume if that sub-portion was not included in data received from another storage volume; instead, the cited portions of Takeda describe a process performed to initialize a volume.

In response to these arguments, the Examiner states that paragraphs 61-62 of Takeda teach “issuing a data read request to the primary host.” Final Office Action mailed November 6, 2007 (“FOA2”), p. 9. The Appellants note that this is not consistent with the actual teachings of Takeda, which state that the “secondary disk array device 200B issues a data read request to the primary disk array device 200A.” Accordingly, the data read request is issued directly between the disk array devices, not to the primary host.

The Examiner also states that “the journal copy is noted as being the request portion that was read and the initial copy being the non-transferred data, which was not included in the requested portion and is therefore being read from the PVOL.” FOA2, p. 10. The Appellants again note that the source of the data copied by both the journal copy and the initial copy is Takeda’s primary volume. Neither operation reads data from

Takeda's secondary volume; instead, both operations copy data from Takeda's primary volume to Takeda's secondary volume. Accordingly, neither operation teaches or suggests "when a sub-portion of the portion of the data is not included in the requested portion [received from the second host, which was requested to provide the requested portion of a copy of the data in a second data storage], reading the sub-portion from the first data storage."

Finally, in the cited portion of Takeda, all data is copied from the primary volume to the secondary volume. Nothing in the cited portions of Takeda teaches or suggests that a read is performed by reading data from the primary volume and from requested data received from a second host that can access the requested data on a secondary volume, or vice versa. Additionally, nothing in the cited portions of Takeda teaches or suggests reading from the secondary volume for any reason, let alone to perform a read.

Claim 1 is patentable of the cited art for the foregoing reasons, as are its dependent claims 2-10. Claims 13-21 are patentable over the cited art for similar reasons and are thus grouped with claim 1 for purposes of this appeal.

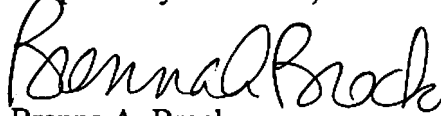
II. Claims 11-12 are patentable over DeKoning in view of Takeda and Carlson

Claims 11 and 12 stand rejected under 35 U.S.C. §103(a) as being unpatentable over DeKoning in view of Takeda and further in view of Carlson. Appellants respectfully request the withdrawal of this rejection for at least the foregoing reasons set forth with respect to claim 1, and thus these claims are also grouped with claim 1 for purposes of this appeal.

CONCLUSION

The appellants respectfully submit that claims 1-21 are allowable over the cited references for at least the above-stated reasons. The appellants respectfully request that the Board reverse the rejections of these claims.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "Brenna A. Brock". The signature is fluid and cursive, with the first name "Brenna" being more prominent than the last name "Brock".

Brenna A. Brock
Attorney for Appellant
Reg. No. 48,509
512-439-5087
512-439-5099 (fax)

VIII. CLAIMS APPENDIX

1. (Previously Presented) A method comprising:
receiving a request to read a portion of data from first data storage, wherein the request is received by a receiving module of a first host, the first host can access the first data storage, and the first host cannot access second data storage;
requesting a requested portion of a copy of the data in the second data storage from a second host that can access the second data storage;
receiving the requested portion from the second host; and
reading the portion of the data by
reading the requested portion received from the second host, and
when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage.
2. (Original) The method of claim 1 further comprising:
determining that a second portion of the data in the first data storage is unavailable;
creating a third data storage upon performing the determining, wherein the first host can access the third data storage; and
causing each subsequent change to the data in the first data storage to be written to the third data storage.
3. (Previously Presented) The method of claim 2 wherein
when the portion of the data comprises an updated portion in the third data storage, the reading the portion of the data comprises reading the updated portion from the third data storage.

4. (Original) The method of claim 2 wherein the second portion of the data is unavailable because the second portion of the data is corrupted.
5. (Original) The method of claim 2 wherein the second portion of the data is unavailable because a device of the first data storage is unavailable.
6. (Original) The method of claim 2 further comprising:
replicating data in the third data storage to fourth data storage accessible by the second host, wherein the fourth data storage cannot be accessed by the first host.
7. (Original) The method of claim 1 wherein the copy of the data in the second data storage was copied from a previous version of the data in the first data storage at a previous point in time.
8. (Original) The method of claim 1 wherein the data in the second data storage is a log of changes made to data in the first data storage after a previous point in time; and
the requested portion is a set of changes in the log of changes, wherein each change in the set of changes comprises a change to the portion of the data, wherein the change was made after the previous point in time.
9. (Original) The method of claim 1 wherein the requesting the requested portion comprises:
identifying a set of changed regions of a first plurality of regions of the first data storage using a set of indicators, wherein each indicator of the set indicates whether at least one change was made to data in a respective region of the first data storage, and;
adding each region of the set of changed regions to the requested portion.

10. (Previously Presented) The method of claim 9 further comprising:
determining whether the data in each region of the first plurality of regions of the first data storage is synchronized with the copy of the data in a corresponding region of a second plurality of regions of the second data storage; and
when the data in one region of the first plurality of regions is not synchronized with the copy of the data in the corresponding region of the second plurality of regions,
identifying a set of unsynchronized regions of the first data storage, wherein each region in the set of unsynchronized regions is unsynchronized with a corresponding region of the second data storage, and
forcing replication of the data in the set of unsynchronized regions to the copy of the data in the second data storage prior to requesting the requested portion.

11. (Previously Presented) The method of claim 10 wherein the determining whether the data in each region of the first data storage is synchronized with the copy of the data in the corresponding region of the second data storage comprises
determining whether a lag in replication from the first data storage to the second data storage exists, and
when the lag exists, determining that the first data storage and the second data storage are unsynchronized.

12. (Previously Presented) The method of claim 11 wherein the determining whether the lag in replication from the first data storage to the second data storage exists comprises:
examining a replication map for the first data storage, wherein the replication map comprises an indicator for each region of the first plurality of regions, the indicator for each region indicates whether data in a respective region of the first data storage have changed but have not yet been replicated; and

when at least one respective region of the first plurality of regions has the indicator, determining that the lag exists.

13. (Previously Presented) A system comprising:

first receiving means of a first host for receiving a request to read a portion of data from first data storage, wherein the first host can access the first data storage, and the first host cannot access second data storage;

requesting means for requesting a requested portion of a copy of the data in the second data storage from a second host that can access the second data storage;

second receiving means for receiving the requested portion from the second host; and

reading means for reading the portion of the data by

reading the requested portion received from the second host, and

when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage.

14. (Original) The system of claim 13 further comprising:

determining means for determining that a second portion of the data in the first data storage is unavailable;

creating means for creating a third data storage upon performing the determining, wherein the first host can access the third data storage; and

causing means for causing each subsequent change to the data in the first data storage to be written to the third data storage.

15. (Original) The system of claim 14 further comprising:

second reading means for reading an updated portion from the third data storage if the portion of the data comprises the updated portion.

16. (Previously Presented) A system comprising:
a first receiving module of a first host to receive a request to read a portion of data from first data storage, wherein the first host can access the first data storage, and the first host cannot access second data storage;
a requesting module to request a requested portion of a copy of the data in the second data storage from a second host that can access the second data storage;
a second receiving module to receive the requested portion from the second host;
and
a reading module to read the portion of the data by
reading the requested portion received from the second host, and
when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage.

17. (Original) The system of claim 16 further comprising:
a determining module to determine that a second portion of the data in the first data storage is unavailable;
a creating module to create a third data storage upon performing the determining, wherein the first host can access the third data storage; and
a causing module to cause each subsequent change to the data in the first data storage to be written to the third data storage.

18. (Original) The system of claim 17 further comprising:
a second reading module to read an updated portion from the third data storage if the portion of the data comprises the updated portion.

19. (Previously Presented) A computer-readable medium comprising:
first receiving instructions to receive a request to read a portion of data from first data storage, wherein a first host can access the first data storage, and the first host cannot access second data storage;

requesting instructions to request a requested portion of a copy of the data in the second data storage from a second host that can access the second data storage;
second receiving instructions to receive the requested portion from the second host; and
reading instructions to read the portion of the data by
reading the requested portion received from the second host, and
when a sub-portion of the portion of the data is not included in the requested portion received from the second host, reading the sub-portion from the first data storage.

20. (Previously Presented) The computer-readable medium of claim 19 further comprising:

determining instructions to determine that a second portion of the data in the first data storage is unavailable;
creating instructions to create a third data storage upon performing the determining, wherein the first host can access the third data storage; and
causing instructions to cause each subsequent change to the data in the first data storage to be written to the third data storage.

21. (Original) The computer-readable medium of claim 20 further comprising:
second reading instructions to read an updated portion from the third data storage
if the portion of the data comprises the updated portion.

IX. EVIDENCE APPENDIX

None

X. RELATED PROCEEDINGS APPENDIX

None